

Remarks

By this paper, amendments are made to further specify the meaning of “quick restart” as used in the present application and claims. As noted previously, the word “quick” refers to a short time interval between the failure and restarting of one node. More particularly, the independent claims are amended herein to specify that the *quick restart at the one node occurs prior to detection of the failure and expulsion of the one node from the membership group due to the failure*. Support for this amendment can be found throughout the application as filed. For example, reference the discussion at page 22 of the specification regarding bounced nodes. No new matter is added to the application by any amendment presented herein. Claims 1-4 & 9-19 remain pending.

Claims 1-4 & 9-19 stand rejected under 35 U.S.C. §102(e) as being anticipated by Knop et al. (U.S. Patent No. 6,885,644; hereinafter Knop). This rejection is respectfully, but most strenuously, traversed to any extent deemed applicable to the amended claims presented, and reconsideration and withdrawal thereof are requested.

The remarks presented below have been prepared in consultation with Felipe Knop, the named inventor in the applied patent, as well as the lead inventor of the present application. A Declaration by Felipe Knop in support of one or more of Applicants’ positioned outlined herein can be provided if the Examiner believes it beneficial to advancement of the current application. *Although the amendments submitted are believed to clearly place all claims in condition for allowance, the Examiner is encouraged to telephone Applicants’ undersigned representative should any issue remain unresolved.*

The present invention is an enhancement upon the basic teachings of Knop. This enhancement relates to the occurrence of a *quick restart* at a node in a membership group in a distributed, multi-node data processing system in which nodes communicate liveness indicia in the form of heartbeat signals via adapters coupled to each node.

Applicants recite (for example, in claim 1) a technique which includes:

- subsequent to a *failure and quick restart* at one node of a membership group, receiving a signal from at least one other node of the membership group at the one node experiencing the failure and quick restart, wherein the failure and quick restart deletes locally stored membership group information at the one node, *and wherein the quick restart at the one node occurs prior to detection of the failure and expulsion of the node from the membership group due to the failure.*

As defined in each of the independent claims presented herewith, Applicants address the problem of a *quick restart* at the one node of a membership group. Specifically, the *quick restart* is defined as a restart occurring at the one node prior to detection of the failure at the one node and expulsion of the one node from the membership group due to the failure. A careful reading of Knop, as well as the other art of record, fails to uncover any teaching or suggestion of a technique for specifically addressing quick restarts at a node of a membership group as defined in the claims presented.

Cited against Applicants' use of the term "quick restart" in the Advisory Action is column 8, line 55 – column 9, line 31 of Knop. The inactivity and reactivity of node 2 in this example in Knop is *not* a quick restart, as now defined in Applicants' independent claims. Specifically, the inactivity of node 2 is detected in Knop and node 2 is expelled from the membership group prior to reactivity of node 2. Thus, the inactivity and reactivity of node 2 is not a quick restart within the characterizations of Applicants' independent claims. Applicants define "quick" relative to the timing for node and/or adapter failure detection in a connected network of nodes. That is, a "quick restart" occurs at a node when the node restarts prior to detection of the failure at the node and expulsion of the node from the membership group.

A careful reading of Knop fails to uncover any discussion of a *quick restart, per se*. As expressly defined in Applicants' independent claims, the current invention is directed to a method of detecting a situation in which a liveness daemon running on one or more nodes has been subject to a rapid restart. More particularly, the present invention is directed (in one aspect) to a method for determining the existence of such a quick restart event and for providing proper indication thereof to other nodes within a network, with the objective of voiding group inconsistencies which are situations in which one node set sees another node set fail in some way without the other node set being aware of the fact that the first node set has also failed. (See page 1, lines 8-14 of the specification.)

With the occurrence of a quick restart on one node, Applicants recite that the failure and quick restart *deletes locally stored membership group information at the one node*. A careful reading of Knop fails to uncover any teaching or suggestion of this concept within the recited language wherein the quick restart at the one node occurs prior to detection of the failure and expulsion of the one node from the membership group due to the failure. Not only is there no quick restart described in Knop as specifically defined in the claims presented, but there is no discussion that the stopping at the one node during the quick restart deletes locally stored membership group information. Such information could alternatively have been maintained in non-volatile storage.

Applicants' invention further recites:

- sending to from the one node to at least one other node, a first message, *which includes at least indicia of occurrence of the quick restart at the one node, the sending being responsive to receipt of the signal at the one node*.

In this regard, the Advisory Action alleges:

The Applicants' specification fails to teach that message sent from one node indicates an occurrence of a quick restart. The specification teaches that detection of the quick restart from the one node is determined by the another node and not by a message sent from the one node. When the one node that restarted attempts to join the another node, the another node detects that the one node is already a member of the group and determines that since the one node is attempting to join and has lost its group information, a restart has occurred at the one node, thereby concluding that a quick restart has occurred and then expelling the one node from the group. (See figure 6 of Applicants' specification.) This is different than what is being argued, which is that a first message contains an indication of the quick restart.

The above conclusions stated in the Advisory Action are believed clearly incorrect.

For example, FIG. 6 of Applicants' specification explains the problem addressed by the present invention, but not the present invention. Applicants' independent claims each recite that the first message *includes at least indicia of occurrence of the quick restart at the one node*. This is supported, by way of example, in the discussion at page 22, line 11 – page 23, line 9 of the specification, wherein protocol for sending a new “NOT YOUR NEIGHBOR” message back from the one node experiencing the failure and quick restart to the sender of the heartbeat message is described. This “NOT YOUR NEIGHBOR” message supports Applicants' recited

invention set forth in the independent claims, wherein the one node experiencing the failure and restart sends a message (e.g., a NOT YOUR NEIGHBOR message) which includes at least indicia of occurrence of the quick restart at the one node. By the one node messaging the another node that it is NOT YOUR NEIGHBOR, the one node is providing indicia to the another node that a quick restart has occurred at the one node.

Cited against this element of Applicants' invention is column 5, lines 42-67 of Knop. These lines state:

In order to monitor the health and connectivity of the adapters in each network, all adapters in the network should attempt to form an "adapter membership group" (AMG), which is a group containing all network adapters that can communicate with each other in the network.

Note that each node may belong to several AMGs, one for each of its network adapters.

To determine the set of adapters that are alive in each network, and adapter membership protocol is run in each of the networks.

As explained further below, adapters that are alive form an Adapter Membership Group (AMG), where members are organized (by way of example only) in a virtual ring topology. To ensure that all group members are alive, each periodically sends "HEART BEAT" messages to its "down-stream neighbor" and monitors the "HEART BEAT" messages from its "upstream neighbor". Protocols are run when adapters fail or when new adapters become functional. The goal of such protocols is to guarantee that the membership group contains at each moment all the adapters that can communicate with each other.

Each group has a "Group Leader" (GL) and a "Backup Group Leader". The group lease is responsible for coordinating the group protocols, and the backup group leader is ...

Again, a careful reading of this material fails to uncover any relevancy to Applicants' recited process for handling quick restarts at one node. The cited lines simply relate to a topology services statement process running at each node. The daemon handles certain aspects of topology propagation facilitating the concepts described in the Knop patent. However, Applicants respectfully submit that the concepts claimed in the present application are simply not described or suggested by Knop.

For example, a careful reading of the cited lines fails to uncover any suggestion that the one node having the quick restart, responsive to receiving a signal from at least one other node of the membership group, sends a first message *which includes at least indicia of occurrence of the quick restart at the one node*. The Knop patent fails to discuss *quick restart, per se* (as defined in the claims presented), and the problems associated therewith. Thus, there is no message which originates from the one node which includes indicia of occurrence of a quick restart at the one node. Further, there is no teaching or suggestion in Knop that such a message is sent from the one node responsive to receipt of a signal from the at least one other node.

Yet further, Applicants' technique (as recited in claim 1) also includes:

- determining at the at least one other node, *from the indicia of occurrence of the quick restart and from locally stored membership group information indicating prior membership of the one node in the membership group, the existence of a quick restart at the one node, and responding thereto by sending a second message from the at least one other node to another node of the membership group which indicates that the one node is to be expelled from the membership group.*

The above-cited lines 42-67 of column 5 are again referenced in the final Office Action for a teaching of this concept. However, Applicants respectfully submit that the processing set forth is simply not taught, suggested or inherent in Knop, irrespective of the interpretation given to Applicants' recited occurrence of a *quick restart*.

In the above-recited aspect of Applicants' invention, the at least one other node (i.e., the node originally sending the signal to the one node having the quick restart) determines from the indicia of occurrence of the quick restart and from its locally stored membership group information indicating prior membership of the one node in the membership group, the existence of a quick restart at the one node. Responding to this determination, the at least one other node *sends a second message to another node in the membership group which indicates that the one node is to be expelled from the membership group*. In contrast, the "HEART BEAT" message and monitoring protocol of Knop simply does not address this aspect of the invention.

Note that if the Examiner is analogizing the phrase *quick restart* to the start of a node, then there is no sense in another node detecting its start and sending out a message that it should be expelled from the membership group. The Knop patent is directed to managing inclusion of active nodes in membership groups.

Applicants respectfully submit that their recited invention patentably distinguishes over the teachings, suggestions and implications of Knop. The procedure recited by Applicants in the independent claims presented is directed to the problem of inconsistencies in group membership occurring as a result of a quick restart. This problem could arise where a liveness daemon at a node is stopped and then restarted quickly.

In situations where the liveness daemon running on one of the nodes is stopped and restarted in a short period of time, certain consistency problems can be engendered. For example, typically it happens that when the liveness daemon restarts, for each local adapter, a message is transmitted which “proclaims” the existence and the willingness of the sending node to become a group leader; that is, in generic terms, a request to know which other nodes are “out there”. However, the other nodes in the group still consider the restarting node (and/or adapter) as being a part of the previous group. Accordingly, group membership is no longer consistent in the sense that there is a lack of symmetry among the various nodes with regards to the “known” status of the other nodes. When the situation is caused by the “quick restart” of the liveness daemon, it is referred to in the application as the “bouncing node” problem or scenario.

This is the problem to which the recited invention is directed. Thus, the phrase *quick restart* has particular meaning in the present application as now expressly defined in the independent claims. A careful reading of Knop fails to uncover any discussion of addressing either of the above-noted problems related to *quick restart* of a node.

Additionally, a careful reading of Knop fails to uncover any discussion or implication that responsive to a quick restart, locally stored membership group information at the one node is deleted. No discussion is provided in the final Office Action why such a step necessarily flows from the teachings of Knop.

Further, Applicants recite in each independent claim that the one other node determines the existence of the quick restart at the one node from the indicia of occurrence of the quick restart and from locally stored membership group information indicating prior membership of the one node in the membership group. There is simply no discussion in Knop of a procedure for ascertaining the existence of a quick restart at one node, let alone the particular procedure recited by Applicants in the independent claims.

Yet further, Applicants recite in response to determining that there was a quick restart at one node, the one other node sends a second message to another node of the membership group which indicates that the one node is to be *expelled* from the membership group. This detection and transmission of a message indicating that one node is to be expelled from a membership group is clearly distinct from the teachings of Knop. The restarted node can communicate with other nodes in the membership group, but for the reasons noted above, the restart must be identified and the restarted node initially expelled from the membership group to avoid the inconsistencies noted. Subsequent to the expulsion, the restarted node can then rejoin the membership group, e.g., using the protocols disclosed by Knop. Advantageously, Applicants' recited facility causes the one node experiencing the quick restart to be expelled from the previous membership group sooner than possible in the prior art. No similar facility is taught or suggested by Knop.

It is well settled that there is no anticipation of a claim unless a single prior art reference discloses: (1) all the same elements of the claimed invention; (2) found in the same situation as the claimed invention; (3) united in the same way as the claimed invention; and (4) in order to perform the identical function as the claimed invention. In this instance, Knop fails to disclose various aspects of Applicants' invention as recited in the independent claims presented, and as a result, does not anticipate (or even render obvious) Applicants' invention. A careful reading of Knop fails to uncover any discussion of the occurrence of a *quick restart* at one node of a membership group. Thus, the problems addressed by Applicants' recited invention are simply not present in the teachings and suggestions of Knop. Additionally, Applicants' independent claims recite a particular process for handling the quick restart to ensure that the quickly restarted node is expelled from the previous membership group expeditiously.

For at least the above reasons, Applicants respectfully submit that independent claims 1, 4 & 14 patentably distinguish over the teachings of Knop. Reconsideration and withdrawal of the rejection based thereon is therefore requested.

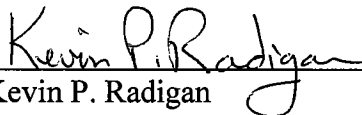
The dependent claims are believed patentable for the same reasons as the independent claims from which they directly or ultimately depend, as well as for their own additional characterizations. For example, claims 3, 6, 10, 11, 16 & 17 further characterize the quick restart indicia. Certain of these claims specify that the indicia comprises each of the three listed components. A careful reading of Knop fails to uncover any teaching or suggestion of sending an indication of a difference in instantiation number for the one node's adapter ID listed in the adapter membership group. Column 6, lines 10-67 of Knop are cited in this regard. However, this material does not discuss an indication of a difference in instantiation number of a quickly restarted node for the node's adapter ID listed in the adapter membership group. This process is simply not part of the Knop patent protocol.

Additionally, claims 8, 13 & 19 further specify the signal, first message and second message, and thereby the protocol for the fast notification of prior nodes of the membership group of the occurrence of the quick restart at the one node. A careful reading of Knop fails to uncover any teaching, suggestion or implication of a message comprising a "NOT YOUR NEIGHBOR" message. The "NOT YOUR NEIGHBOR" message was specifically written by Applicant Felipe Knop, subsequent to filing of the applied Knop patent. The *NOT YOUR NEIGHBOR* message is a message used by the one node having the quick restart to tell other nodes that it quickly restarted.

The application is believed to be in condition for allowance, and such action is respectfully requested.

If a telephone conference would be of assistance in advancing prosecution of this application, Applicants' undersigned attorney invites the Examiner to telephone him at the number provided.

Respectfully submitted,



Kevin P. Radigan
Attorney for Applicants
Registration No.: 31,789

Dated: October 17, 2006.

HESLIN ROTHENBERG FARLEY & MESITI P.C.
5 Columbia Circle
Albany, New York 12203-5160
Telephone: (518) 452-5600
Facsimile: (518) 452-5579